Freeform Search

Database:	US Pre-Grant Publication Full-Text Database US Patents Full-Text Database US OCR Full-Text Database EPO Abstracts Database JPO Abstracts Database Derwent World Patents Index IBM Technical Disclosure Bulletins				
Term:	L17 and 15				
Display: 10 Documents in <u>Display Format</u> : - Starting with Number 1 Generate: O Hit List O Hit Count O Side by Side O Image					
	Search				
Search History					

DATE: Saturday, May 05, 2007 Purge Queries Printable Copy Create Case

Set Name side by side	Query	<u>Hit</u> Count	Set Name result set
DB=P	GPB, USPT, USOC, EPAB, JPAB, DWPI, TDBD; PLUR=YES; OP=OR		
<u>L18</u>	L17 and 15	1	<u>L18</u>
<u>L17</u>	(asynchronously with streams with (parallel or concurrent\$2 or simultaneously))	54	<u>L17</u>
L16	L15 and l14	4	<u>L16</u>
<u>L15</u>	(((log\$1 or block) with file) with (parallel or simultaneously or concurrent\$2))	1296	<u>L15</u>
<u>L14</u>	L13 and 15	32	<u>L14</u>
<u>L13</u>	(asynchronous with (parallel or concurrent\$3 or simultaneous\$2))	6970	<u>L13</u>
<u>L12</u>	11 and (707/8).ccls.	1	<u>L12</u>
<u>L11</u>	(707/204).ccls. and 11	. 6	<u>L11</u>
<u>L10</u>	(recover\$3 with backup) and 11	10	<u>L10</u>
<u>L9</u>	L8 and 15	3	<u>L9</u>
<u>L8</u>	(archiv\$3 with log\$1 with (parallel or simultaneously or concurrent\$2))	32	<u>L8</u>
<u>L7</u>	7177866	3	<u>L7</u>
<u>L6</u>	L5 and 11	3	<u>L6</u>

<u>L5</u>	(asynchronous with (backup or archiv\$3))	361	<u>L5</u>
<u>L4</u>	L3	0	<u>L4</u>
<u>L3</u>	11 and L2	0	<u>L3</u>
<u>L2</u>	(asynchronous with (backup or archiv\$3)).ab.	36	<u>L2</u>
<u>L1</u>	(asynchronous with (parallel or concurrent\$3 or simultaneous\$2) with (log or block\$1))	277	<u>L1</u>

END OF SEARCH HISTORY

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1	"752607".apn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR .	OFF	2007/05/05 08:24
L2	36	(asynchronous with (backup or archiv\$3)).ab.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/05/05 09:38
L3	15	(asynchronous with (parallel or concurrent\$3 or simultaneous\$2) with log)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/05/05 08:35
L4	0	2 and 3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/05/05 08:32
L5	355	(asynchronous with (backup or archiv\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/05/05 08:32
L6	3	3 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/05/05 08:32
L7	0	(asynchronous with (parallel or concurrent\$3 or simultaneous\$2) with logs)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/05/05 08:36

Sign in

Google

 Web
 Images
 Video
 News
 Maps
 more »

 asynchronous backup simultaneously archive
 Search
 Advanced Search Preferences

Web Results 1 - 10 of about 188,000 for asynchronous backup simultaneously archive log. (0.23 seconds)

Evaluation of remote backup algorithms for transaction-processing ... Subhash Bhalla, Stuart E. Madnick, Asynchronous Backup and Initialization ... If the redo-log is applied continuously as if the backup site were part of a ... portal.acm.org/citation.cfm?id=185836&coll=portal&dl=ACM - Similar pages

Replication Group - Cluster Schema - GigaSpaces Online Help ...

Backup space does not need to process redo log data. Backup space needs to ... Logs the operation and the relevant Entry UID into an asynchronous redo log. ...

www.gigaspaces.com/wiki/display/GS/Replication+Group+-+Cluster+Schema - 64k - Cached - Similar pages

Database Recovery

If asynchronous I/O is supported, the recovery time may only be slightly reduced ... The restored backup files will have earlier log sequence numbers than ... www.lsbu.ac.uk/oracle/oracle7/server/doc/SCN73/ch24.htm - 37k - Cached - Similar pages

Managing Archived Redo Logs

A database **backup**, together with online and archived redo **log** files, ... to one of the **log** files, the ARCn processes can **simultaneously archive** up to three ... www.lc.leidenuniv.nl/awcourse/oracle/server.920/a96521/archredo.htm - 88k - Cached - Similar pages

Command Reference

The next database **backup** will prune the expired entries from the history file. ... that the user exit function cannot be called to store **archive log** files. ... publib.boulder.ibm.com/infocenter/db2v7luw/topic/com.ibm.db2v7.doc/db2n0/sqln0046.htm - 27k - <u>Cached</u> - <u>Similar pages</u>

<u>DBAzine.com:</u> An Oracle Instructor's Guide To Oracle Data Guard Administrators are able to specify synchronous or **asynchronous** network ... The Fetch **Archive Log** Process (FAL) is a background Oracle process that runs on ... www.dbazine.com/oracle/or-articles/foot7 - 38k - Cached - Similar pages

Data Guard vs. HADR

For resolution of large **archive-log** gaps, in order to minimize the delay in ... In contrast to Oracle, DB2 offers a basic set of **backup** and recovery ... www.oracle.com/technology/deploy/availability/htdocs/DataGuardHADR.html - 79k - Cached - Similar pages

[PDF] Oracle database remote replication using Sun StorEdge Data ... File Format: PDF/Adobe Acrobat - View as HTML online logs, archive logs, and control file be mirrored. Administrators must place the database into Hot Backup mode so that all of the data can be ... www.sun.com/storagetek/white-papers/oracledatabase_remotereplication.pdf Similar pages

Database Recovery

A database **backup** consists of backups of the physical files (all ... If ARCn encounters an error when attempting to **archive** a **log** group (for example, ... www.csee.umbc.edu/help/oracle8/server.815/a67781/c28recov.htm - 86k -

Cached - Similar pages

NDMP.org - Protocol Specification Summary

It is also possible to use the NDMP to **simultaneously backup** to two tape drives ... These messages allows the NDMP server to make entries in the **backup log**. ... www.ndmp.org/info/spec_summary.shtml - 18k - <u>Cached</u> - <u>Similar pages</u>

Google Groups results for asynchronous backup simultaneously archive log

PITR Functional Design v2 for 7.5 - comp.databases.postgresql.hack ... - Mar 9, 2004 E Class :: Oracle-----Day 2 - Mithun's Group - Nov 14, 2005 Why do people switch to Linux? - comp.os.linux.advocacy - Oct 31, 2005

Result Page: 1 <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u> <u>Next</u>

Try Google Desktop: search your computer as easily as you search the web.

asynchronous backup simultaneous | Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2007 Google

Sign in

<u>Google</u>

 Web
 Images
 Video
 News
 Maps
 more »

 asynchronous backup corrupted simultaneous
 Search
 Advanced Search Preferences

Web Results 1 - 10 of about 66,200 for asynchronous backup corrupted simultaneously archive logs. (0.29)

Data Guard vs. HADR

For resolution of large **archive-log** gaps, in order to minimize the delay in ... In contrast to Oracle, DB2 offers a basic set of **backup** and recovery ... www.oracle.com/technology/deploy/availability/htdocs/DataGuardHADR.html - 79k - Cached - Similar pages

[PDF] Oracle Data Guard

File Format: PDF/Adobe Acrobat automatically detected by Data Guard and the necessary **archive logs** are ... all the protection modes (i.e. whether using synchronous or **asynchronous** mode ... www.oracle.com/technology/deploy/availability/pdf/40056_Ray_doc.pdf - <u>Similar pages</u>

Managing Archived Redo Logs

A database **backup**, together with online and archived redo **log** files, ... to one of the **log** files, the ARCn processes can **simultaneously archive** up to three ... www.lc.leidenuniv.nl/awcourse/oracle/server.920/a96521/archredo.htm - 88k - Cached - Similar pages

<u>Using Oracle8 -- Chapter 14 - Performing Database Recovery</u>

Do you have all **archive logs** accessible without any break in the sequence numbers? ... the redo **logs** are misplaced, removed, **corrupt**, or from an old **backup**. ... docs.rinet.ru/O8/ch14/ch14.htm - 53k - <u>Cached</u> - <u>Similar pages</u>

Database Recovery

The restored **backup** files will have earlier **log** sequence numbers than those of the ... online redo **log**) might become **corrupt** due to a system failure. ... www.lsbu.ac.uk/oracle/oracle7/server/doc/SCN73/ch24.htm - 37k - <u>Cached</u> - <u>Similar pages</u>

Database Recovery

information about backups of datafiles and **archive logs** ... Incorrect behavior by Recovery Manager cannot **corrupt** the database. ... www.csee.umbc.edu/help/oracle8/server.815/a67781/c28recov.htm - 86k - Cached - Similar pages

[РРТ] HP standard dark template

File Format: Microsoft Powerpoint - <u>View as HTML</u> hot database files, such as DBF files and **archive log** files which should not be ... **Corrupt** Block detection. supported. requires Begin/End. **Backup** stmts. ... h21007.www2.hp.com/dspp/files/unprotected/database/HP3KOracle.ppt - Similar pages

Oracle Undocumented Parameters

There is a new init.ora parameter, LOG_BLOCKS_DURING_BACKUP, that controls whether block images ar written to the redo **log** during hot **backup**. ... www.fortunecity.com/skyscraper/oracle/699/orahtml/misc/undoc.html - 27k - Cached - Similar pages

Technical . DB2AdministrationOverview

When you **back up** a table space online, it is still available for use, and **simultaneous** updates are recorded in the **logs**. When you perform an online restore ... faemalia.org/wiki/view/Technical/DB2AdministrationOverview - 36k - Cached - Similar pages

Contents

BACKUP... DELETE INPUT With **Corrupt Archive Logs ...** Identifying Bottlenecks with Synchronous I/O · Identifying Bottlenecks with **Asynchronous** I/O ... docs.huihoo.com/oracle/docs/B19306_01/backup.102/b14191/toc.htm - 96k - Cached - Similar pages

Google Groups results for asynchronous backup corrupted simultaneously archive logs



50 Page Study Guide - alt.certification.network-plus - Aug 23, 2000 Update 2.03 available - comp.unix.unixware.misc - Nov 21, 1995 The NeXT-FAQ available! (Part II) - de.comp.sys.next - Jan 25, 1996

Result Page:

1 2 3 4 5 6 7 8 9 10

Next

Try Google Desktop: search your computer as easily as you search the web.

asynchronous backup corrupted sim

Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2007 Google



Home | Login | Logout | Access Information | Alerts |

Search Results BROWSE SEARCH IEEE XPLORE GUIDE Results for "(asynchronous backup corrupted simultaneously archive logs<in>metadata)" ⊠e-mail Your search matched 0 documents. A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order. » Search Options View Session History **Modify Search** New Search (asynchronous backup corrupted simultaneously archive logs<in>metadata) «Search Check to search only within this results set » Key Indicates full text access IEEE Journal or **IEEE JNL** Magazine No results were found. **IET JNL** IET Journal or Magazine Please edit your search criteria and try again. Refer to the Help pages if you need assistan IEEE Conference **IEEE CNF** Proceeding **IET Conference IET CNF** Proceeding

Indexed by inspec'

IEEE STD IEEE Standard

Help Contact Us Privacy &: © Copyright 2006 IEEE -



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library O The Guide

asynchronous backup corrupted simultaneously archive logs



THE ACCIDICITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used asynchronous backup corrupted simultaneously archive logs

Found 4,833 of 200,192

Sort results by

relevance

Save results to a Binder

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form

Open results in a new window

Results 1 - 20 of 200 Best 200 shown

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

Relevance scale

1 Peer-to-peer infrastructure: Pastiche: making backup cheap and easy

Landon P. Cox, Christopher D. Murray, Brian D. Noble

December 2002 ACM SIGOPS Operating Systems Review, Volume 36 Issue SI

Publisher: ACM Press

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings

Backup is cumbersome and expensive. Individual users almost never back up their data, and backup is a significant cost in large organizations. This paper presents Pastiche, a simple and inexpensive backup system. Pastiche exploits excess disk capacity to perform peer-to-peer backup with no administrative costs. Each node minimizes storage overhead by selecting peers that share a significant amount of data. It is easy for common installations to find suitable peers, and peers with high ove ...

2 ARIES: a transaction recovery method supporting fine-granularity locking and partial





rollbacks using write-ahead logging

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz March 1992 ACM Transactions on Database Systems (TODS), Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(5.23 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

High speed on-line backup when using logical log operations



David B. Lomet

May 2000 ACM SIGMOD Record, Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00, Volume 29 Issue 2

Publisher: ACM Press

Full text available: 📆 <u>pdf(220.69 KB)</u> Additional Information: <u>full citation, abstract, references, citings, index</u> terms

Media recovery protects a database from failures of the stable medium by maintaining an extra copy of the database, called the backup, and a media recovery log. When a failure occurs, the database is "restored" from the backup, and the media recovery log is used to roll forward the database to the desired time, usually the current time. Backup must be both fast and "on-line", i.e. concurrent with on-going update activity. Conventional online backup sequentially copies ...

4 Disaster recovery techniques for database systems



Manhoi Choy, Hong Va Leong, Man Hon Wong November 2000 Communications of the ACM

Publisher: ACM Press

Full text available: pdf(412.04 KB) Additional Information: full citation, references, index terms

5 OceanStore: an architecture for global-scale persistent storage



Systems Review, Proceedings of the ninth international conference on Architectural support for programming languages and operating systems ASPLOS-IX, Volume 28, 34 Issue 5, 5

Publisher: ACM Press

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(166.53 KB) terms

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers, data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks; monitoring also enhances performance through pro-active movement ...

6 OceanStore: an architecture for global-scale persistent storage



John Kubiatowicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishan Gummadi, Sean Rhea, Hakim Weatherspoon, Westley Weimer, Chris Wells, Ben

November 2000 ACM SIGPLAN Notices, Volume 35 Issue 11

Publisher: ACM Press

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(1.47 MB) terms

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers, data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks; monitoring also enhances performance through pro-active movement ...

7 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data



consistency between database and file systems

Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on

Management of data SIGMOD '02

Publisher: ACM Press

Full text available: Topdf(1.44 MB)

Additional Information: full citation, abstract, references, index terms

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

8 RPC-V: Toward Fault-Tolerant RPC for Internet Connected Desktop Grids with Volatile Nodes

Samir Djilali, Thomas Herault, Oleg Lodygensky, Tangui Morlier, Gilles Fedak, Franck

November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing SC

Publisher: IEEE Computer Society

Full text available: pdf(230.75 KB) Additional Information: full citation, abstract

RPC is one of the programming models envisioned for the Grid. In Internet connected Large Scale Grids such as Desktop Grids, nodes and networks failures are not rare events. This paper provides several contributions, examining the feasibility and limits of fault-tolerant RPC on these platforms. First, we characterize these Grids from their fundamental features and demonstrate that their applications scope should be safely restricted to stateless services. Second, we present a new fault-tolerant ...

9 Practical byzantine fault tolerance and proactive recovery

Min
 No
 No

Miguel Castro, Barbara Liskov

November 2002 ACM Transactions on Computer Systems (TOCS), Volume 20 Issue 4

Publisher: ACM Press

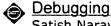
Full text available: pdf(1.63 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>, <u>review</u>

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

Keywords: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

10 BugNet: Continuously Recording Program Execution for Deterministic Replay



Satish Narayanasamy, Gilles Pokam, Brad Calder

May 2005 ACM SIGARCH Computer Architecture News, Proceedings of the 32nd Annual International Symposium on Computer Architecture ISCA '05, Volume 33 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(204.07 KB) Additional Information: full citation, abstract, citings, index terms

Significant time is spent by companies trying to reproduce and fix the bugs that occur for

released code. To assist developers, we propose the BugNet architecture to continuously record information on production runs. The information collected before the crash of a program can be used by the developers working in their execution environment to deterministically replay the last several million instructions executed before the crash. BugNet is based on the insight that recording the register file ...

11 The Alpine file system



M. R. Brown, K. N. Kolling, E. A. Taft

November 1985 ACM Transactions on Computer Systems (TOCS), Volume 3 Issue 4

Publisher: ACM Press

Full text available: pdf(2.95 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Alpine is a file system that supports atomic transactions and is designed to operate as a service on a computer network. Alpine's primary purpose is to store files that represent databases. An important secondary goal is to store ordinary files representing documents, program modules, and the like. Unlike other file servers described in the literature, Alpine uses a log-based technique to implement atomic file update. Another unusual aspect of Alpine is that it performs all commu ...

12 Distributed logging for transaction processing



Dean S. Daniels, Alfred Z. Spector, Dean S. Thompson

December 1987 ACM SIGMOD Record, Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87, Volume

16 Issue 3 **Publisher: ACM Press**

Full text available: pdf(1.51 MB)

Additional Information: full citation, abstract, references, citings, index terms

Increased interest in using workstations and small processors for distributed transaction processing raises the question of how to implement the logs needed for transaction recovery. Although logs can be implemented with data written to duplexed disks on each processing node, this paper argues there are advantages if log data is written to multiple log server nodes. A simple analysis of expected logging loads leads to the conclusion that a high performance, microprocessor b ...

13 Data integration and sharing II: Scientific data repositories: designing for a moving



target

Etzard Stolte, Christoph von Praun, Gustavo Alonso, Thomas Gross

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(739.27 KB) terms

Managing scientific data warehouses requires constant adaptations to cope with changes in processing algorithms, computing environments, database schemas, and usage patterns. We have faced this challenge in the RHESSI Experimental Data Center (HEDC), a datacenter for the RHESSI NASA spacecraft. In this paper we describe our experience in developing HEDC and discuss in detail the design choices made. To successfully accommodate typical adaptations encountered in scientific data management systems ...

14 The Recovery Manager of the System R Database Manager



Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 ACM Computing Surveys (CSUR), Volume 13 Issue 2

Publisher: ACM Press

Full text available: Tpdf(1.75 MB) Additional Information: full citation, references, citings, index terms

15 Beating the I/O bottleneck: a case for log-structured file systems

John Ousterhout, Fred Douglis

January 1989 ACM SIGOPS Operating Systems Review, Volume 23 Issue 1

Publisher: ACM Press

Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, citings, index terms

CPU speeds are improving at a dramatic rate, while disk speeds are not. This technology shift suggests that many engineering and office applications may become so I/O-limited that they cannot benefit from further CPU improvements. This paper discusses several techniques for improving I/O performance, including caches, battery-backed-up caches, and cache logging. We then examine in particular detail an approach called log-structured file systems, where the file system's only representation ...

16 The Conquest file system: Better performance through a disk/persistent-RAM hybrid



design

An-I Andy Wang, Geoff Kuenning, Peter Reiher, Gerald Popek August 2006 ACM Transactions on Storage (TOS), Volume 2 Issue 3

Publisher: ACM Press

Additional Information: full citation, abstract, references, index terms Full text available: pdf(1.34 MB)

Modern file systems assume the use of disk, a system-wide performance bottleneck for over a decade. Current disk caching and RAM file systems either impose high overhead to access memory content or fail to provide mechanisms to achieve data persistence across reboots. The Conquest file system is based on the observation that memory is becoming inexpensive, which enables all file system services to be delivered from memory, except for providing large storage capacity. Unlike caching, Con ...

Keywords: Persistent RAM, file systems, performance measurement, storage management

17 The process group approach to reliable distributed computing



Kenneth P. Birman

December 1993 Communications of the ACM, Volume 36 Issue 12

Publisher: ACM Press

Full text available: pdf(6.00 MB)

Additional Information: full citation, references, citings, index terms

Keywords: fault-tolerant process groups, message ordering, multicast communication

18 A Self-Organizing Storage Cluster for Parallel Data-Intensive Applications Hong Tang, Aziz Gulbeden, Jingyu Zhou, William Strathearn, Tao Yang, Lingkun Chu November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing SC '04

Publisher: IEEE Computer Society

Full text available: pdf(330.26 KB) Additional Information: full citation, abstract

Cluster-based storage systems are popular for data-intensive applications and it is desirable yet challenging to provide incremental expansion and high availability while achieving scalability and strong consistency. This paper presents the design and implementation of a self-organizing storage cluster called Sorrento, which targets dataintensive workload with highly parallel requests and low write-sharing patterns. Sorrento automatically adapts to storage node joins and departures, and the sys ...

19 The Zebra striped network file system

John H. Hartman, John K. Ousterhout

December 1993 ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP

'93, Volume 27 Issue 5

Publisher: ACM Press

Full text available: pdf(1.93 MB)

Additional Information: full citation, abstract, references, citings, index terms

Zebra is a network file system that increases throughput by striping file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this increase ...

20 Rx: treating bugs as allergies---a safe method to survive software failures

Feng Qin, Joseph Tucek, Jagadeesan Sundaresan, Yuanyuan Zhou

October 2005 ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05, Volume 39 Issue

Publisher: ACM Press

Full text available: pdf(245.29 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u>

Many applications demand availability. Unfortunately, software failures greatly reduce system availability. Prior work on surviving software failures suffers from one or more of the following limitations: Required application restructuring, inability to address deterministic software bugs, unsafe speculation on program execution, and long recovery time. This paper proposes an innovative safe technique, called Rx, which can quickly recover programs from many types of software bugs, both det ...

Keywords: availability, bug, reliability, software failure

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player

Sign in

Google

 Web
 Images
 Video
 News
 Maps
 more »

 background synchronization archive logs simul
 Search
 Advanced Search Preferences

Web Results 1 - 10 of about 133,000 for background synchronization archive logs simultaneously asynchi

Log Apply Services

In this mode, **log** transport services **archive logs** to the standby site, and **log** apply ... processes are able to apply the archived redo **logs simultaneously** ... lc.leidenuniv.nl/awcourse/oracle/server.920/a96653/standbyreco.htm - 133k - Cached - Similar pages

[DOC] Oracle added 36 new packages 9i

File Format: Microsoft Word - View as HTML

These logs keep track of the changes to the master tables and are used by the materialized

views during refresh synchronization. ...

www.tcoug.org/Archive/Winter2003/9i%20Packages.doc - Similar pages

PITR Functional Design v2 for 7.5

This is an **asynchronous** call, so postgresql will not wait for that xlog to complete archiving. This call is currently not likely to be called **simultaneously** ... **archives**.postgresql.org/pgsql-hackers/2004-03/msg00295.php - 34k - Cached - Similar pages

[PDF] 40144 - Data Guard SQL Apply - Back to the Future!

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> updated, they can be used **simultaneously** for other tasks such as reporting, ... database, even last Sunday's as long as you have the **Archive Logs** available ... www.oracle.com/technology/deploy/availability/pdf/40144_Carpenter_doc.pdf - <u>Similar pages</u>

[PDF] Oracle7 Server

File Format: PDF/Adobe Acrobat - View as HTML

ADT has implemented an **Archive Log** Repository on the LAN and are sending. redo to that site using the **Log** Writer **asynchronous** mode (LGWR ASYNC) which ... www.oracle.com/technology/deploy/availability/pdf/1173_Carpenter_WP.pdf - <u>Similar pages</u> [<u>More results from www.oracle.com</u>]

Internet Archive Search: mediatype: "software" AND subject: "Internet"

This Web log file tool analyzes logs generated by IIS, Apache and other Web ... asynchronous messaging system access and XML remote procedure calls to Web ... www.archive.org/search.php?query=mediatype%3A%22software%22%20AND% 20subject%3A%22Internet%22&page=12 - 86k - Cached - Similar pages

Asynchronous coordinated commit replication and dual write with ...

(b) The tp subsystem usually provides automatic redo **log** retention/management. ... **Synchronization** may be synchronous or **asynchronous**. ... www.freepatentsonline.com/7177866.html - 194k - Cached - Similar pages

Method of processing a synchronization point in a database ...

Method of processing a **synchronization** point in a database management system to assure a database version using update **logs** from accumulated transactions ... www.freepatentsonline.com/5778388.html - 92k - <u>Cached</u> - <u>Similar pages</u> [<u>More results from www.freepatentsonline.com</u>]

[PDF] Scientific Data Repositories - Designing for a Moving Target

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> remote **archive** times out, no results are available); query re-. sults are not cached, and there is no data **synchronization** be-... acm.org/sigmod/sigmod03/eproceedings/papers/r13p03.pdf - Similar pages

(PPT) Introduction

File Format: Microsoft Powerpoint - <u>View as HTML</u>
Importance of Clocks & **Synchronization**. Avoiding **simultaneous** access of ... an entry into their undo and redo **logs**; Nodes send agreement/abort messages ... www.cs.uga.edu/~laks/ADCS-Materials/Introduction.ppt - <u>Similar pages</u>

Result Page: 1 2 3 4 5 6 7 8 9 10 Next

Download Google Pack: free essential software for your PC

background synchronization archive Search

<u>Search within results</u> | <u>Language Tools</u> | <u>Search Tips</u> | <u>Dissatisfied? Help us improve</u>

Google Home - Advertising Programs - Business Solutions - About Google

©2007 Google

Sign in

Google

 Web
 Images
 Video
 News
 Maps
 more »

 background synchronization archive logs simul
 Search
 Advanced Search Preferences

Web Results 1 - 10 of about 133,000 for background synchronization archive logs simultaneously asynchi

Log Apply Services

In this mode, **log** transport services **archive logs** to the standby site, and **log** apply ... processes are able to apply the archived redo **logs simultaneously**. ... lc.leidenuniv.nl/awcourse/oracle/server.920/a96653/standbyreco.htm - 133k - Cached - Similar pages

[DOC] Oracle added 36 new packages 9i

File Format: Microsoft Word - View as HTML

These logs keep track of the changes to the master tables and are used by the materialized

views during refresh synchronization. ...

www.tcoug.org/Archive/Winter2003/9i%20Packages.doc - Similar pages

PITR Functional Design v2 for 7.5

This is an **asynchronous** call, so postgresql will not wait for that xlog to complete archiving. This call is currently not likely to be called **simultaneously** ... **archive**s.postgresql.org/pgsql-hackers/2004-03/msg00295.php - 34k - Cached - Similar pages

[PDF] 40144 - Data Guard SQL Apply - Back to the Future!

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> updated, they can be used **simultaneously** for other tasks such as reporting, ... database, even last Sunday's as long as you have the **Archive Logs** available ... www.oracle.com/technology/deploy/availability/pdf/40144_Carpenter_doc.pdf - <u>Similar pages</u>

[PDF] Oracle7 Server

File Format: PDF/Adobe Acrobat - View as HTML

ADT has implemented an **Archive Log** Repository on the LAN and are sending. redo to that site using the **Log** Writer **asynchronous** mode (LGWR ASYNC) which ... www.oracle.com/technology/deploy/availability/pdf/1173_Carpenter_WP.pdf - <u>Similar pages</u> [More results from www.oracle.com]

Internet Archive Search: mediatype: "software" AND subject: "Internet"

This Web **log** file tool analyzes **logs** generated by IIS, Apache and other Web ... **asynchronous** messaging system access and XML remote procedure calls to Web ... www.archive.org/search.php?query=mediatype%3A%22software%22%20AND% 20subject%3A%22Internet%22&page=12 - 86k - <u>Cached</u> - <u>Similar pages</u>

Asynchronous coordinated commit replication and dual write with ...

(b) The tp subsystem usually provides automatic redo log retention/management. ... Synchronization may be synchronous or asynchronous. ... www.freepatentsonline.com/7177866.html - 194k - Cached - Similar pages

Method of processing a synchronization point in a database ...

Method of processing a **synchronization** point in a database management system to assure a database version using update **logs** from accumulated transactions ... www.freepatentsonline.com/5778388.html - 92k - <u>Cached</u> - <u>Similar pages</u> [More results from www.freepatentsonline.com]

[PDF] Scientific Data Repositories - Designing for a Moving Target

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> remote **archive** times out, no results are available); query re-. sults are not cached, and there is no data **synchronization** be- ... acm.org/sigmod/sigmod03/eproceedings/papers/r13p03.pdf - <u>Similar pages</u>

[PPT] Introduction

File Format: Microsoft Powerpoint - <u>View as HTML</u>
Importance of Clocks & **Synchronization**. Avoiding **simultaneous** access of ... an entry into their undo and redo **logs**; Nodes send agreement/abort messages ... www.cs.uga.edu/~laks/ADCS-Materials/Introduction.ppt - <u>Similar pages</u>

Result Page: 1 2 3 4 5 6 7 8 9 10 Next

Download Google Pack: free essential software for your PC

background synchronization archive | Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2007 Google



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library C The Guide

background synchronization archive logs asynchronous simulta



HE ACTIONS IN LIBRARY

Feedback Report a problem Satisfaction

Terms used background synchronization archive logs asynchronous simultaneously

window

Found 13,054 of 200,192

Sort results by

Display

results

relevance

expanded form

Save results to a Binder ? Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

Best 200 shown

Relevance scale

1 ARIES: a transaction recovery method supporting fine-granularity locking and partial



rollbacks using write-ahead logging

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz March 1992 ACM Transactions on Database Systems (TODS), Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(5.23 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

FAB: building distributed enterprise disk arrays from commodity components



Yasushi Saito, Svend Frølund, Alistair Veitch, Arif Merchant, Susan Spence October 2004 ACM SIGARCH Computer Architecture News, ACM SIGOPS Operating Systems Review, ACM SIGPLAN Notices, Proceedings of the 11th

international conference on Architectural support for programming languages and operating systems ASPLOS-XI, Volume 32 , 38 , 39 Issue 5 , 5 , 11

Publisher: ACM Press

Full text available: ndf(671.67 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

This paper describes the design, implementation, and evaluation of a Federated Array of Bricks (FAB), a distributed disk array that provides the reliability of traditional enterprise arrays with lower cost and better scalability. FAB is built from a collection of bricks, small storage appliances containing commodity disks, CPU, NVRAM, and network interface cards. FAB deploys a new majority-voting-based algorithm to replicate or erasure-code logical blocks across bricks and a reconfigurati ...

Keywords: consensus, disk array, erasure coding, replication, storage, voting

3 ISIS: an adaptive, trilingual conversational system with interleaving interaction and



delegation dialogs

Helen Meng, P. C. Ching, Shuk Fong Chan, Yee Fong Wong, Cheong Chat Chan September 2004 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 11 Issue 3

Publisher: ACM Press

Full text available: pdf(3.71 MB) Additional Information: full citation, abstract, references, index terms

ISIS (Intelligent Speech for Information Systems) is a trilingual spoken dialog system (SDS) for the stocks domain. It handles two dialects of Chinese (Cantonese and Putonghua) as well as English---the predominant languages in our region. The system supports spoken language queries regarding stock market information and simulated personal portfolios. The conversational interface is augmented with a screen display that can capture mouse-clicks as well as textual input by typing or stylus-writing. ...

Keywords: Human-computer spoken language interface, interaction and delegation dialogs

4 Data integration and sharing II: Scientific data repositories: designing for a moving



target

Etzard Stolte, Christoph von Praun, Gustavo Alonso, Thomas Gross

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Full text available: pdf(739.27 KB)

Additional Information: full citation, abstract, references, citings, index terms

Managing scientific data warehouses requires constant adaptations to cope with changes in processing algorithms, computing environments, database schemas, and usage patterns. We have faced this challenge in the RHESSI Experimental Data Center (HEDC), a datacenter for the RHESSI NASA spacecraft. In this paper we describe our experience in developing HEDC and discuss in detail the design choices made. To successfully accommodate typical adaptations encountered in scientific data management systems ...

5 Synchronization and recovery in a client-server storage system



E. Panagos, A. Biliris

August 1997 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 6 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(205.25 KB) Additional Information: full citation, abstract, citings, index terms

Client-server object-oriented database management systems differ significantly from traditional centralized systems in terms of their architecture and the applications they target. In this paper, we present the client-server architecture of the EOS storage manager and we describe the concurrency control and recovery mechanisms it employs. EOS offers a semi-optimistic locking scheme based on the multi-granularity two-version two-phase locking protocol. Under this scheme, multiple concurrent reade ...

Keywords: Checkpoint, Client-server architecture, Object management, Concurrency control, Locking, Logging, Recovery, Transaction management

6 Collocation and virtual collocation: Benefits of synchronous collaboration support for



an application-centered analysis team working on complex problems: a case study
John M. Linebarger, Andrew J. Scholand, Mark A. Ehlen, Michael J. Procopio

November 2005 Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work GROUP '05

Publisher: ACM Press

Full text available: pdf(682.03 KB) Additional Information: full citation, abstract, references, index terms

A month-long quasi-experiment was conducted using a distributed team responsible for modeling, simulation, and analysis. Six experiments of three different time durations (short, medium, and long) were performed. The primary goal was to discover if synchronous collaboration capability through a particular application improved the ability of the team to form a common mental model of the analysis problem(s) and solution(s). The results indicated that such collaboration capability did improve the f ...

Keywords: benefits of collaboration, collaboration experiments, collaboration frameworks, common mental models, synchronous collaboration

7 The Escrow transactional method

Patrick E. O'Neil

December 1986 ACM Transactions on Database Systems (TODS), Volume 11 Issue 4

Publisher: ACM Press

Full text available: pdf(2.36 MB)

Additional Information: full

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>, <u>review</u>

A method is presented for permitting record updates by long-lived transactions without forbidding simultaneous access by other users to records modified. Earlier methods presented separately by Gawlick and Reuter are comparable but concentrate on "hotspot" situations, where even short transactions cannot lock frequently accessed fields without causing bottlenecks. The Escrow Method offered here is designed to support nonblocking record updates by transactions that are "lo ...

8 A method for implementing lock-free shared-data structures

Greg Barnes

August 1993 Proceedings of the fifth annual ACM symposium on Parallel algorithms and architectures SPAA '93

Publisher: ACM Press

Full text available: 🔂 pdf(978.61 KB) Additional Information: full citation, references, citings, index terms

⁹ Garbage collection for a client-server persistent object store

Laurent Amsaleg, Michael J. Franklin, Olivier Gruber

August 1999 ACM Transactions on Computer Systems (TOCS), Volume 17 Issue 3

Publisher: ACM Press

Full text available: pdf(267.18 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

We describe an efficient server-based algorithm for garbage collecting persistent object stores in a client-server environmnet. The algorithm is incremental and runs concurrently with client transactions. Unlike previous algorithms, it does not hold any transactional locks on data and does non require callbacks to clients. It is fault-tolerant, but performs very little logging. The algorithm has been designed to be integrated into existing systems, and therefore it works with standard i ...

Keywords: client-server system, logging, persistent object-store, recovery

10
Activity-based computing: support for mobility and collaboration in ubiquitous

computing

E. Bardram

September 2005 Personal and Ubiquitous Computing, Volume 9 Issue 5

Publisher: Springer-Verlag

Full text available: pdf(412.31 KB) Additional Information: full citation, abstract, citings

This paper presents the design philosophy of activity-based computing (ABC), which addresses mobility and cooperation in human work activities. Furthermore, it presents the ABC framework, which is a ubiquitous computing infrastructure supporting ABC. The idea of ABC and the aim of the ABC framework is to: (1) support human activity by managing its collection of work tasks on a computer, (2) support mobility by distributing activities across heterogeneous computing environments, (3) suppor ...

Keywords: Activity-based computing, Computer supported cooperative work, Contextaware computing, Pervasive healthcare, State management, Ubiquitous computing

11 The SNAP-1 parallel AI prototype

R. F. DeMara, D. I. Moldovan

April 1991 ACM SIGARCH Computer Architecture News, Proceedings of the 18th annual international symposium on Computer architecture ISCA '91, Volume

19 Issue 3
Publisher: ACM Press

Full text available: pdf(1.07 MB)

Additional Information: full citation, references, citings, index terms

12 Special issue: Al in engineering

D. Sriram, R. Joobbani

April 1985 ACM SIGART Bulletin, Issue 92

Publisher: ACM Press

Full text available: pdf(8.79 MB) Additional Information: full citation, abstract

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

13 <u>Designing presentations for on-demand viewing</u>

Liwei He, Jonathan Grudin, Anoop Gupta

December 2000 Proceedings of the 2000 ACM conference on Computer supported cooperative work CSCW '00

Publisher: ACM Press

Full text available: pdf(487.54 KB)

Additional Information: full citation, abstract, references, citings, index terms

Increasingly often, presentations are given before a live audience, while simultaneously being viewed remotely and recorded for subsequent viewing on-demand over the Web. How should video presentations be designed for web access? How is video accessed and used online? Does optimal design for live and on-demand audiences conflict? We examined detailed behavior patterns of more than 9000 on-demand users of a large corpus of professionally prepared presentations. We find that as many people ac ...

Keywords: digital library, streaming media, video on-demand

¹⁴ Fault-scalable Byzantine fault-tolerant services



Michael Abd-El-Malek, Gregory R. Ganger, Garth R. Goodson, Michael K. Reiter, Jay J. Wylie October 2005 ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05, Volume 39 Issue-

Publisher: ACM Press

Full text available: The pdf(356.66 KB) Additional Information: full citation, abstract, references, index terms

A fault-scalable service can be configured to tolerate increasing numbers of faults without significant decreases in performance. The Query/Update (Q/U) protocol is a new tool that enables construction of fault-scalable Byzantine fault-tolerant services. The optimistic quorum-based nature of the Q/U protocol allows it to provide better throughput and faultscalability than replicated state machines using agreement-based protocols. A prototype service built using the Q/U protocol outperfor ...

Keywords: byzantine fault-tolerance, fault-scalability, quorums, replicated state machines, services

15 A structural view of the Cedar programming environment



Daniel C. Swinehart, Polle T. Zellweger, Richard J. Beach, Robert B. Hagmann

August 1986 ACM Transactions on Programming Languages and Systems (TOPLAS),

Volume 8 Issue 4

Publisher: ACM Press

Full text available: pdf(6.32 MB)

Additional Information: full citation, abstract, references, citings, index

This paper presents an overview of the Cedar programming environment, focusing on its overall structure—that is, the major components of Cedar and the way they are organized. Cedar supports the development of programs written in a single programming language, also called Cedar. Its primary purpose is to increase the productivity of programmers whose activities include experimental programming and the development of prototype software systems for a high-performance personal computer. T ...

16 RPC-V: Toward Fault-Tolerant RPC for Internet Connected Desktop Grids with Volatile Nodes



Samir Djilali, Thomas Herault, Oleg Lodygensky, Tangui Morlier, Gilles Fedak, Franck Cappello

November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing SC '04

Publisher: IEEE Computer Society

Full text available: pdf(230.75 KB) Additional Information: full citation, abstract

RPC is one of the programming models envisioned for the Grid. In Internet connected Large Scale Grids such as Desktop Grids, nodes and networks failures are not rare events. This paper provides several contributions, examining the feasibility and limits of faulttolerant RPC on these platforms. First, we characterize these Grids from their fundamental features and demonstrate that their applications scope should be safely restricted to stateless services. Second, we present a new fault-tolerant ...

17 Fast detection of communication patterns in distributed executions



November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research CASCON '97

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the

execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

18 OAI in action: Repository synchronization in the OAI framework

Xiaoming Liu, Kurt Maly, Mohammad Zubair, Michael L. Nelson

May 2003 Proceedings of the 3rd ACM/IEEE-CS joint conference on Digital libraries JCDL '03

Publisher: IEEE Computer Society

Full text available: pdf(98.38 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u>

The Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) began as an alternative to distributed searching of scholarly eprint repositories. The model embraced by the OAI-PMH is that of metadata harvesting, where value-added services (by a "service provider") are constructed on cached copies of the metadata extracted from the repositories of the harvester's choosing. While this model dispenses with the well known problems of distributed searching, it introduces the problem of synch ...

19 Nomadic radio: speech and audio interaction for contextual messaging in nomadic



<u>environments</u>

Nitin Sawhney, Chris Schmandt

September 2000 ACM Transactions on Computer-Human Interaction (TOCHI), Volume 7 Issue 3

Publisher: ACM Press

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: The pdf(648.76 KB) terms.

Mobile workers need seamless access to communication and information services while on the move. However, current solutions overwhelm users with intrusive interfaces and ambiguous notifications. This article discusses the interaction techniques developed for Nomadic Radio, a wearable computing platform for managing voice and text-based messages in a nomadic environment. Nomadic Radio employs an auditory user interface, which synchronizes speech recognition, speech synthesis, nonspeech audio ...

Keywords: adaptive interfaces, contextual interfaces, interruptions, nonspeech audio, notifications, passive awareness, spatial listening, speech interaction, wearable computing

20 ReVive: cost-effective architectural support for rollback recovery in shared-memory





multiprocessors

Milos Prvulovic, Zheng Zhang, Josep Torrellas

May 2002 ACM SIGARCH Computer Architecture News, Proceedings of the 29th annual international symposium on Computer architecture ISCA '02, Proceedings of the 29th annual international symposium on Computer architecture ISCA '02, Volume 30 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: Additional Information: full citation, abstract, references, citings, index terms Publisher Site

This paper presents ReVive, a novel general-purpose rollback recovery mechanism for shared-memory multiprocessors. ReVive carefully balances the conflicting requirements of availability, performance, and hardware cost. ReVive performs checkpointing, logging, and distributed parity protection, all memory-based. It enables recovery from a wide class of errors, including the permanent loss of an entire node. To maintain high performance,

ReVive includes specialized hardware that performs frequent o ...

Keywords: fault tolerance, shared-memory multiprocessors, rollback recovery, recovery, BER, logging, parity, checkpointing, availability

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player